Scammers Allegedly Used an AI-Cloned Lawyer's Voice to Defraud Italian Private Bank Fideuram

In February 2026, scammers used AI-generated voice cloning to impersonate a prominent lawyer, convincing the chairman of the Italian private bank Fideuram to authorize fraudulent wire transfers. The attackers, posing as the CEO of the parent company Intesa Sanpaolo via WhatsApp and a cloned voice call, successfully directed approximately €195 million to overseas accounts. While authorities and the bank managed to recover a significant portion of the funds, approximately €36 million remains missing after being converted into cryptocurrencies.

Description

Scammers allegedly used an AI-cloned lawyer's voice to validate a WhatsApp request impersonating the CEO of Italian bank Intesa Sanpaolo. Fideuram, its private-banking subsidiary, reportedly transferred about €195 million overseas after its chairman accepted the instructions. Part of the money was recovered, but tens of millions reportedly remained missing. Milan prosecutors opened a fraud investigation.

AI Incident Database ↗ Incident date: 2026-02-23

Risk Classification

Based on MIT AI Risk Repository Causal and Domain Taxonomies and the EU AI Act Risk Tiers.

Risk Domain / Subdomain

The incident involved scammers using AI voice cloning to impersonate a trusted lawyer to execute a high-value financial fraud against a bank executive.

Domain	Subdomain
4 - Malicious actors	4.3 - Fraud, scams, and targeted manipulation

Causal Classification

The incident was initiated by human scammers who intentionally deployed a voice cloning AI tool post-deployment to commit financial fraud.

Entity	Intent	Timing
Human	Intentional	Post-deployment

EU AI Act Risk Tier

EU AI Act Risk Tier	EU AI Act Risk Classification Reasoning
3 - Limited Risk	Limited Risk: The system used is an AI-generated content tool (deepfake/voice clone). The EU AI Act classifies AI-generated content such as deepfakes under Limited Risk, requiring transparency obligations, though in this case, it was used maliciously by bad actors.

AI System

Developer	Model family	Model name
Unidentified AI system	Unidentified AI system	—

Behaviour type

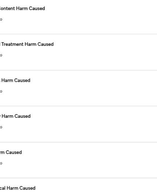
Tool

Harm Severity

Harm caused by the incident is broken out into:

- Explicitly reported **directly caused** harm (e.g. financial loss from an AI-enabled fraud)
- Explicitly reported **indirectly caused** harm (e.g. financial value of reputational damage following an incident)
- **Inferred** harm, that was not explicitly described in the reports but was inferred by the model and may be worthy of further investigation.

Severity ratings are based on this scale. **Highest Direct Severity Score** is the highest Direct Severity for this incident in any single category.



Highest Direct Severity Score (in any single category)

Substantial

Physical Harm Caused	Physical Severity Direct	Physical Severity Indirect
Reported: No	Negligible	Negligible

Infrastructure Harm Caused	Infrastructure Severity Direct	Infrastructure Severity Indirect
Reported: No	Negligible	Negligible

Property Harm Caused	Property Severity Direct	Property Severity Indirect
Reported: No	Negligible	Negligible

Financial Harm Caused	Financial Severity Direct	Financial Severity Indirect
Reported: The reports explicitly describe a financial loss of 95 million euros, of which 36 million euros remains missing. This occurred in 1 primary incident, making the average loss per occurrence 95 million euros initially, with a net permanent loss of 36 million euros. Directly caused: The fraudulent transaction directly caused Fideuram to transfer 95 million euros to overseas accounts, with 36 million euros currently unrecovered. Inferred additional harm: The bank and its parent company likely incurred additional indirect financial costs related to legal fees, internal investigations, and security audits following the incident.	Substantial	Negligible

Environmental Harm Caused	Environmental Severity Direct	Environmental Severity Indirect
Reported: No	Negligible	Negligible

Malicious Content Harm Caused	Malicious Content Severity Direct	Malicious Content Severity Indirect
Reported: No	Negligible	Negligible

Differential Treatment Harm Caused	Differential Treatment Severity Direct	Differential Treatment Severity Indirect
Reported: No	Negligible	Negligible

Civil Rights Harm Caused	Civil Rights Severity Direct	Civil Rights Severity Indirect
Reported: No	Negligible	Negligible

Democracy Harm Caused	Democracy Severity Direct	Democracy Severity Indirect
Reported: No	Negligible	Negligible

Privacy Harm Caused	Privacy Severity Direct	Privacy Severity Indirect
Reported: No	Negligible	Negligible

Psychological Harm Caused	Psychological Severity Direct	Psychological Severity Indirect
Reported: No Inferred additional harm: It is likely that Tommaso Mulosini experienced significant stress, anxiety, and reputational distress leading to his resignation, though this is not explicitly detailed as clinical psychological harm.	Negligible	Negligible

Epistemic Harm Caused	Epistemic Severity Direct	Epistemic Severity Indirect
Reported: Yes, the report describes the fabrication of a lawyer's voice to deceive a bank executive. Directly caused: The scammers used AI voice cloning to fabricate the voice of a prominent lawyer, misleading the bank chairman into believing the transaction instructions were legitimate.	Negligible	Negligible

CSEA Harm Caused	CSEA Severity Direct	CSEA Severity Indirect
Reported: No	Negligible	Negligible

Numbers of People Affected

These numbers are produced by the LLM tool from the incident reports. Treat as indicative and verify against the underlying sources.

Number of occurrences reported

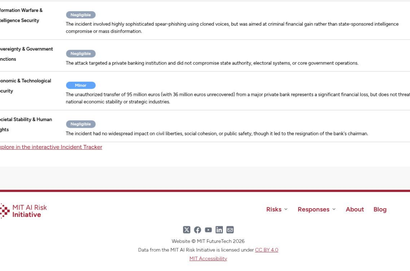
1

Total people reportedly exposed to risk	Total people reportedly harmed
2	1

Total people inferred exposed to risk	Total people inferred harmed
5	1

Potential Causes of Incident

Fishbone / Ishikawa Diagram — potential causes for the incident identified by the LLM analysis, organised by category. Treat as indicative.



Alleged Entities

As identified by the AI Incident Database.

Alleged deployer	Alleged developer
Voice cloning impersonation scammers, Synthetic media creators, Scammers targeting Fideuram, Scammers, Impersonation scammers, Deepfake creators	Voice cloning technology developers, Synthetic media generation technology developers, Synthetic audio generation technology developers, Deepfake technology developers

Alleged harmed parties	
Victims of voice cloning impersonation scams, Victims of impersonation scams, Privacy, Paolo Nastasi, Paolo Molesini, Lawyers, Intesa Sanpaolo, Financial scam victims, Fideuram, Epistemic integrity, Banks	—

Information Quality

Confidence	Missing Information
High	No

Reason for Confidence	Missing Information Detail
The reports from multiple reputable sources (ANSA, Corriere della Sera, Reuters) are highly consistent regarding the core facts: the date, the target (Fideuram/Molesini), the amount stolen (95 million euros), the amount recovered, and the use of AI voice cloning to impersonate a lawyer. There is very little ambiguity about the sequence of events.	N/A

Ambiguities Identified	
Minor discrepancies exist regarding the exact amount recovered and currently missing (some reports mention 39.5 million missing, others 36 million), but the overall scale is clear.	—

Alternative interpretations	
None. It is clearly a case of AI-enabled financial fraud.	—

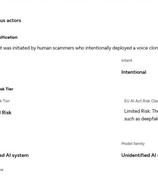
National Security Assessment

The National Security Impact Assessment is based on this framework

The assessment uses voluntarily submitted incident reports and automated classification by an LLM. It should therefore be treated as indicative and all sources/analysis must be checked before being used or referenced.

National Security Impact Summary

International scammers used AI voice cloning to impersonate a lawyer, defrauding an Italian private bank of 95 million euros (36 million euros unrecovered). While a highly sophisticated financial crime representing an evolved AI threat capability, the incident has minor national security implications as it did not target critical infrastructure, government functions, or economic stability.



Overall National Security Impact Score

Based on the highest individual category rating

Minor

Affected Nations / Actors

Scope of Impact	Primary Target
Multiple nations	Italy

Alleged Perpetrator	Other Affected Actors
Unknown	China, Portugal

Threat Characteristics

Imminence Level — Long-term

The active crisis has concluded, but the incident represents an ongoing strategic concern regarding AI-enabled financial fraud capabilities.

Autonomy Level — Human-controlled

The voice cloning AI was used as a tool directly controlled and deployed by human fraudsters to execute the social engineering scheme.

Novelty Level — Evolved capability

While voice cloning fraud is an established threat, the scale of the financial target and the sophistication of the multi-stage social engineering represent an evolved capability.

National Security Impact

Scale: Negligible · Minor · Substantial · Severe · Critical

Physical Security & Critical Infrastructure — Negligible

No physical assets, critical infrastructure, or human safety systems were targeted or affected by this AI-enabled financial fraud.

Information Warfare & Intelligence Security — Negligible

The incident involved highly sophisticated spear-phishing using cloned voices, but was aimed at criminal financial gain rather than state-sponsored intelligence compromise or mass disinformation.

Sovereignty & Government Functions — Negligible

The attack targeted a private banking institution and did not compromise state authority, electoral systems, or core government operations.

Economic & Technological Security — Minor

The unauthorized transfer of 95 million euros (with 36 million euros unrecovered) from a major private bank represents a significant financial loss, but does not threaten national economic stability or strategic industries.

Societal Stability & Human Rights — Negligible

The incident had no widespread impact on civil liberties, social cohesion, or public safety, though it led to the resignation of the bank's chairman.

Explore in the interactive Incident Tracker

Website © MIT FutureTech 2026
Data from the MIT AI Risk Initiative is licensed under CC BY 4.0
MIT Accessibility